DATABASE INDEX VALIDATION MECHANISM

ABSTRACT OF THE DISCLOSURE

A method evaluates a plurality of candidate index sets for a workload of database statements in a database system by first generating baseline statistics for each statement in the workload. An index superset is formed by combining an existing or current index set and a proposed index set. A candidate index set is derived from the index superset, the candidate index being one of the plurality of candidate index sets. Statistics for a statement are generated by first creating an execution plan which represents an efficient series of steps for executing the statement given the candidate index set. The execution plan is evaluated, and statistics based on the evaluation of the execution plan are generated and recorded. The cost of the execution plan is then determined and statistics are generated. Statistics for each candidate index set are rolled up and presented to a user or an index tuning mechanism.